

A direct solver with $O(N)$ complexity for integral equations on one-dimensional domains

A. Gillman, P. Young, P.G. Martinsson

Abstract: An algorithm for the direct inversion of the linear systems arising from Nyström discretization of integral equations on one-dimensional domains is described. The method typically has $O(N)$ complexity when applied to boundary integral equations (BIEs) in the plane with non-oscillatory kernels such as those associated with the Laplace and Stokes' equations. The scaling coefficient suppressed by the “big-O” notation depends logarithmically on the requested accuracy. The method can also be applied to BIEs with oscillatory kernels such as those associated with the Helmholtz and Maxwell equations; it is efficient at long and intermediate wave-lengths, but will eventually become prohibitively slow as the wave-length decreases. To achieve linear complexity, rank deficiencies in the off-diagonal blocks of the coefficient matrix are exploited. The technique is conceptually related to the \mathcal{H} - and \mathcal{H}^2 -matrix arithmetic of Hackbusch and co-workers, and is closely related to previous work on *Hierarchically Semi-Separable* matrices.

1. INTRODUCTION

1.1. Problem formulation. The paper describes techniques for numerically solving equations of the type

$$(1.1) \quad a(t) q(t) + \int_0^T b(t, t') q(t') dt' = f(t), \quad t \in [0, T],$$

where $I = [0, T]$ is an interval on the line, and where $a : I \rightarrow \mathbb{R}$ and $b : I \times I \rightarrow \mathbb{R}$ are given functions. We observe that a boundary integral equation (BIE) such as

$$(1.2) \quad a(\mathbf{x}) q(\mathbf{x}) + \int_{\Gamma} b(\mathbf{x}, \mathbf{x}') q(\mathbf{x}') dl(\mathbf{x}') = f(\mathbf{x}), \quad \mathbf{x} \in \Gamma,$$

where Γ is a simple curve in the plane takes the form (1.1) upon parameterization of the curve. The case of a domain Γ that consists of several non-connected simple curves can also be handled.

Upon discretization, equation (1.1) takes the form

$$(1.3) \quad \mathbf{A} \mathbf{q} = \mathbf{f}$$

where \mathbf{A} is a dense matrix of size, say, $N \times N$. When N is large, standard practice for rapidly solving a system such as (1.3) is to use an iterative solver (such as GMRES, conjugate gradients, etc.) in which the matrix-vector multiplications are accelerated via a “fast” method such as the Fast Multipole Method (FMM) [13], panel clustering [15], Barnes-Hut [2], etc. When the integral equation (1.1) is a Fredholm equation of the second kind, the iteration typically converges rapidly, and a linear solver of effectively $O(N)$ complexity results. In contrast, this papers reviews and extends a number of recently developed *direct* solvers that in a single pass compute a data-sparse representation of a matrix \mathbf{S} (a “solution operator”) that satisfies

$$\mathbf{S} \approx \mathbf{A}^{-1}.$$

Once a representation of \mathbf{S} is available, the solution of (1.3) is of course easily constructed:

$$(1.4) \quad \mathbf{q} \approx \mathbf{S} \mathbf{f}.$$

We will demonstrate that in many important environments (such as, e.g., the BIEs associated with Laplace’s equation in the plane), the matrix \mathbf{S} can be constructed in $O(N)$ operations.

1.2. Applications. The direct solver presented is applicable to most boundary integral equations associated with the classical boundary value problems of mathematical physics (Laplace, elasticity, Helmholtz, Yukawa, Stokes, etc.) with the two important exceptions that it is not efficient for (1) problems involving highly oscillatory kernels such as Helmholtz equation at short wavelengths, and (2) domain boundaries that tend to “fill space” in the sense illustrated in Figure 1.1. We will demonstrate that high accuracy and speed can be maintained even for non-smooth boundaries.

The direct solver is also applicable to many integral equations of the form (1.1) that arise in the analysis of special functions [28], in evaluating conformal maps [24], and in the analysis of two-point boundary value problems [27].

1.3. Advantages of direct solvers. Direct solvers offer several advantages over iterative ones:

Speed-up by large factors for problems involving multiple right hand sides: In many situations, an equation such as (1.1) needs to be solved for several different data functions f . Iterative techniques can only to a limited extent take advantage of the fact that the operator is the same in each solve. For a direct method, on the other hand, each solve beyond the first simply involves applying a pre-computed inverse to a vector. The time required for applying the (compressed) inverse to a vector is typically much smaller than even the time required for a single application of the original operator using standard techniques.

The ability to solve relatively ill-conditioned problems: Direct solvers allow for the rapid and accurate solution of linear systems involving relatively ill-conditioned matrices. In the context of boundary value problems, such ill-conditioning can be caused by physical ill-conditioning (as observed, e.g., when solving the equations of elasticity on domains with high aspect ratios, or when solving scattering problems near a resonant frequency), but may also be introduced as a side-effect of the mathematical formulation (e.g. when a formulation based on a Fredholm equation of the first kind is used, or when a model of a scattering problem introduces so called “spurious” resonances).

Increased reliability: Direct methods are inherently more robust than iterative methods. This point is less important in an academic setting where there is often time to tweak a code until it performs well for a particular problem (for instance by designing a customized pre-conditioner). However, it has proven difficult to design industrial codes using iterative methods and commercial software developers sometimes tend to shun iterative methods in favor of direct ones, even at significant cost in terms of speed.

1.4. How the direct solver works. Letting ε denote a user specified computational tolerance, the direct solver for (1.1) can be viewed as consisting of four steps:

(i) *Quadrature nodes and quadrature weights for a Nyström discretization are created:* The interval $[0, T]$ is split into panels, and Gaussian nodes are placed on each panel. Customized quadrature weights are constructed using the method of [29] which ensures high accuracy even in the presence of weakly singular kernels (and for BIEs on domains with corners).




FIGURE 1.1. Contours for which the direct solver will *not* achieve $O(N)$ complexity.

(ii) *Construction of the coefficient matrix:* The matrix \mathbf{A} in (1.3) is an $N \times N$ matrix that is dense, but whose off-diagonal blocks are to high accuracy rank-deficient. We exploit this fact, and compute an approximant $\mathbf{A}_{\text{approx}}$ which is stored in the data-sparse format very similar to the *Hierarchically Semi-Separable (HSS)* format of [26, 6].

(iii) *Inversion of the coefficient matrix:* The approximant $\mathbf{A}_{\text{approx}}$ of the coefficient matrix is inverted using a variation of the technique of [27, 21] to produce the solution operator $\mathbf{S} = \mathbf{A}_{\text{approx}}^{-1}$. The inversion is exact up to round-off errors.

(iv) *Application of the approximate inverse:* The solution operator \mathbf{S} is applied to the given data vector \mathbf{f} to produce the solution \mathbf{q} , cf. (1.4).

Each of the four steps typically requires $O(N)$ work when applied to the standard equations of mathematical physics (with the two exceptions mentioned in Section 1.2). The constants of proportionality depend on the specific environment, but in general, Step 2 is the most expensive. The cost of Step (iv) is tiny, meaning that the proposed procedure is particularly effective in environments where a sequence of equations with the same coefficient matrix need to be solved.

Remark 1.1. The computations in steps 3 and 4 are independent of the specific problem being solved, and can be implemented as “black-box” codes.

1.5. Relationship to earlier work. The general idea of exploiting rank-deficiencies in the off-diagonal blocks of the matrix \mathbf{A} in (1.3) underlies several “fast” algorithms (e.g. the Fast Multipole Method [13], panel clustering [15], Barnes-Hut [2]) for executing the matrix-vector multiply in an iterative solver. The observation that such rank-deficiencies can be also used in a systematic way to execute matrix inversion, matrix-matrix multiplies, matrix factorizations, etc., was made in the early work on \mathcal{H} -matrices by Hackbusch and co-workers [16]. The basic version of these methods have $O(N(\log N)^p)$ complexity for some small integer p . Certain operations were later accelerated to $O(N)$ complexity in the context of \mathcal{H}^2 -matrices, see [4] and the references therein.

More specifically, the direct solver described in this paper is an evolution of the scheme of [21], which in turn draws on the earlier work [3, 23, 27]. Since [21] appeared, the algorithms have been significantly improved, primarily in that the compression and inversion steps have been separated. In addition to making the presentation much clearer, this separation leads to several concrete improvement, including:

Improved versatility: Separating the task of compression from the task of inversion makes it much easier to apply the direct solver to new applications. If a BIE with a different kernel is to be solved, a slight modification of the compression step (Step (ii) in Section 1.4) is sufficient. It also opens up the possibility of combining the direct solver with generic compression techniques based on randomized sampling, e.g., those described in [20].

Improved quadratures: The version of the algorithm described in this paper is compatible with the quadratures of [17, 5] which enable the handling of BIEs defined on domains with corners, and the quadratures of [29] which simplify the handling of singular kernels.

Improved theoretical analysis: The direct solver is in the present paper expressed transparently as a telescoping matrix factorization. This allows for a simplified error and stability analysis, as illustrated by, e.g., Lemma 3.1 and Corollary 3.2.

Improved interoperability with other data-sparse matrix formats: The new version of the algorithm makes it clear that the data-sparse format used to represent both the coefficient matrix and its inverse are essentially identical to the *Hierarchically Semi-Separable (HSS)* format of

[26, 6]. This opens up the possibility of combining the compression techniques described in this paper with recently developed inversion and factorization algorithms for HSS matrices [8].

Remark 1.2. The paper uses the terms “block separable” (BS) and “hierarchically block separable” (HBS). The HBS format is essentially identical to the HSS format. The terms BS and HBS were introduced for local purposes only since they clarify the description of the algorithm. There is no intention to replace the well-established term “HSS.”

1.6. Outline. Section 2 introduces notation and reviews the Nyström discretization method for integral equations. Section 3 describes an accelerated direct solver based on a simplistic tessellation of an $N \times N$ matrix A into $p \times p$ blocks in such a way that all off-diagonal blocks are rank deficient. This method has complexity $O(p^{-2} N^3 + p^3 k^3)$ where k is the rank of the off-diagonal blocks. To attain better asymptotic complexity, a more complicated hierarchical tessellation of the matrix must be implemented. This data structure is described in Section 4, and an $O(N k^2)$ inversion technique is then described in Section 5. Section 6 describes efficient techniques for computing the data-sparse representation in the first place. Section 7 describes some numerical experiments, and Section 8 describes possible extensions of the work.

2. PRELIMINARIES

This section introduces notation, and briefly reviews some known techniques.

2.1. Notation. We say that a matrix U is *orthonormal* if its columns form an orthonormal set. An orthonormal matrix U preserves geometry in the sense that $|Ux| = |x|$ for every vector x . We use the notation of [10] to denote submatrices: If A is an $m \times n$ matrix with entries $A(i, j)$, and if $I = [i_1, i_2, \dots, i_p]$ and $J = [j_1, j_2, \dots, j_q]$ are two index vectors, then the associated $p \times q$ submatrix is expressed as

$$A(I, J) = \begin{bmatrix} a_{i_1, j_1} & \cdots & a_{i_1, j_q} \\ \vdots & & \vdots \\ a_{i_p, j_1} & \cdots & a_{i_p, j_q} \end{bmatrix}.$$

For column- and row-submatrices, we use the standard abbreviations

$$A(:, J) = A([1, 2, \dots, m], J), \quad \text{and} \quad A(I, :) = A(I, [1, 2, \dots, n]).$$

2.2. The Interpolatory Decomposition (ID). An $m \times n$ matrix B of rank k admits the factorization

$$B = UB(J, :),$$

where $J = [j_1, \dots, j_k]$ is a vector of integers such that $1 \leq j_i \leq m$, and U is a $m \times k$ matrix that contains the $k \times k$ identity matrix I_k (specifically, $U(J, :) = I_k$). Moreover, no entry of U is larger than one. Computing the ID of a matrix is in general combinatorially expensive, but if the restriction on element size of U is relaxed slightly to say that, for instance, each entry of U is bounded by 2, then very efficient schemes are available. See [14, 9] for details.

2.3. Nyström discretization of integral equations in one dimension. In this section, we very briefly describe some variations of the classical Nyström method for discretizing an integral equation such as (1.1). The material is well known and we refer to [1] for details.

For an integral equation with a smooth kernel $k(t, t')$, the Nyström method is particularly simple. The starting point is a quadrature rule for the interval $[0, T]$ with nodes $\{t_i\}_{i=1}^N \subset [0, T]$ and weights $\{\omega_i\}_{i=1}^N$ such that

$$\int_0^T b(t_i, t') q(t') dt' \approx \sum_{j=1}^n b(t_i, t_j) q(t_j) \omega_j, \quad i = 1, 2, \dots, N.$$

Then the discretized version of (1.1) is obtained by enforcing that

$$(2.1) \quad a(t_i) q(t_i) + \sum_{j=1}^n b(t_i, t_j) q(t_j) \omega_j = f(t_i), \quad i = 1, 2, \dots, N.$$

We write (2.1) compactly as

$$\mathbf{A} \mathbf{q} = \mathbf{f},$$

where \mathbf{A} is the $N \times N$ matrix with entries

$$\mathbf{A}(i, j) = \delta_{i,j} a(t_i) + b(t_i, t_j) \omega_j, \quad i, j = 1, 2, 3, \dots, N.$$

where \mathbf{f} is the vector with entries

$$\mathbf{f}(i) = f(t_i), \quad i = 1, 2, 3, \dots, N,$$

and where \mathbf{q} is the approximate solution which satisfies

$$\mathbf{q}(i) \approx q(t_i), \quad i = 1, 2, 3, \dots, N.$$

We have found that using a composite quadrature rule with a 10-point standard Gaussian quadrature on each panel is a versatile and highly accurate choice.

Remark 2.1 (Singular kernels). Some of the numerical examples described in Section 7 involve kernels with logarithmically singular kernels,

$$k(t, t') \sim \log |t - t'|, \quad \text{as } t' \rightarrow t.$$

A standard quadrature rule designed for smooth functions would lose almost all accuracy on the panels where t and t' is close, but this can be remedied by modifying the matrix entries near the diagonal. For instance, when Gaussian quadrature nodes are used, the procedure described in [29] gives very accurate results. Alternatively, the Rokhlin-Kapur [18] procedure starts with a standard trapezoidal rule and modifies the weights near the end points to achieve high order convergence. This is a simpler method than the modified Gaussian rule of [29] but typically also produces lower accuracy.

Remark 2.2 (Contours with corners). Discretizing an integral equation such as (1.2) can be challenging if the contour Γ is not smooth. When $\mathbf{x} \in \Gamma$ is a corner point, the function $\mathbf{x}' \mapsto b(\mathbf{x}, \mathbf{x}')$ typically has a singularity at \mathbf{x} . It has been demonstrated [17, 5] that in many cases of practical interest, it is nevertheless possible to use standard quadrature weights designed for smooth functions, as long as the discretization is locally refined near the corner. The drawback is that such refinement can increase the system size in an undesirable way but as [17] demonstrates, the system size can be reduced via a local pre-computation. In this paper, we demonstrate that it is alternatively possible to use general purpose direct solvers to achieve the same effect.

3. INVERSION OF BLOCK SEPARABLE MATRICES

In this section, we define what it means for a matrix to be “block separable” and describe a simple technique for inverting such a matrix.

Let \mathbf{A} be an $np \times np$ matrix that is blocked into $p \times p$ blocks, each of size $n \times n$:

$$(3.1) \quad \mathbf{A} = \begin{bmatrix} \mathbf{D}_1 & \mathbf{A}_{1,2} & \mathbf{A}_{1,3} & \cdots & \mathbf{A}_{1,p} \\ \mathbf{A}_{2,1} & \mathbf{D}_2 & \mathbf{A}_{2,3} & \cdots & \mathbf{A}_{2,p} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{A}_{p,1} & \mathbf{A}_{p,2} & \mathbf{A}_{p,3} & \cdots & \mathbf{D}_p \end{bmatrix}.$$

We say that \mathbf{A} is “block separable” with “block-rank” k if for $\tau = 1, 2, \dots, p$, there exist $n \times k$ matrices \mathbf{U}_τ and \mathbf{V}_τ such that each off-diagonal block $\mathbf{A}_{\sigma,\tau}$ of \mathbf{A} admits the factorization

$$(3.2) \quad \begin{array}{ccc} \mathbf{A}_{\sigma,\tau} & = & \mathbf{U}_\sigma \quad \tilde{\mathbf{A}}_{\sigma,\tau} \quad \mathbf{V}_\tau^*, \\ n \times n & & n \times k \quad k \times k \quad k \times n \end{array} \quad \sigma, \tau \in \{1, 2, \dots, p\}, \quad \sigma \neq \tau.$$

Observe that the columns of \mathbf{U}_σ must form a basis for the columns of all off-diagonal blocks in row σ , and analogously, the columns of \mathbf{V}_τ must form a basis for the rows in all the off-diagonal blocks in column τ . When (3.2) holds, the matrix \mathbf{A} admits a block factorization

$$(3.3) \quad \begin{array}{ccccc} \mathbf{A} & = & \underline{\mathbf{U}} & \tilde{\mathbf{A}} & \underline{\mathbf{V}}^* \\ np \times np & & np \times kp & kp \times kp & kp \times np \end{array} + \begin{array}{c} \underline{\mathbf{D}} \\ np \times np \end{array},$$

where

$$\mathbf{U} = \text{diag}(\mathbf{U}_1, \mathbf{U}_2, \dots, \mathbf{U}_p), \quad \mathbf{V} = \text{diag}(\mathbf{V}_1, \mathbf{V}_2, \dots, \mathbf{V}_p), \quad \mathbf{D} = \text{diag}(\mathbf{D}_1, \mathbf{D}_2, \dots, \mathbf{D}_p),$$

and

$$\tilde{\mathbf{A}} = \begin{bmatrix} 0 & \tilde{\mathbf{A}}_{12} & \tilde{\mathbf{A}}_{13} & \cdots \\ \tilde{\mathbf{A}}_{21} & 0 & \tilde{\mathbf{A}}_{23} & \cdots \\ \tilde{\mathbf{A}}_{31} & \tilde{\mathbf{A}}_{32} & 0 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}.$$

The block structure of formula (3.3) for $p = 4$ is illustrated below:

$$(3.4) \quad \begin{array}{ccccc} \mathbf{A} & = & \underline{\mathbf{U}} & \tilde{\mathbf{A}} & \underline{\mathbf{V}}^* \\ \begin{array}{|c|c|c|c|} \hline \text{shaded} & \text{shaded} & \text{shaded} & \text{shaded} \\ \hline \text{shaded} & \text{shaded} & \text{shaded} & \text{shaded} \\ \hline \text{shaded} & \text{shaded} & \text{shaded} & \text{shaded} \\ \hline \text{shaded} & \text{shaded} & \text{shaded} & \text{shaded} \\ \hline \end{array} & & \begin{array}{|c|c|c|c|} \hline \text{shaded} & \text{white} & \text{white} & \text{white} \\ \hline \text{white} & \text{shaded} & \text{white} & \text{white} \\ \hline \text{white} & \text{white} & \text{shaded} & \text{white} \\ \hline \text{white} & \text{white} & \text{white} & \text{shaded} \\ \hline \end{array} & & \begin{array}{|c|c|c|c|} \hline \text{shaded} & \text{shaded} & \text{shaded} & \text{shaded} \\ \hline \text{shaded} & \text{shaded} & \text{shaded} & \text{shaded} \\ \hline \text{shaded} & \text{shaded} & \text{shaded} & \text{shaded} \\ \hline \text{shaded} & \text{shaded} & \text{shaded} & \text{shaded} \\ \hline \end{array} & + & \begin{array}{|c|c|c|c|} \hline \text{shaded} & \text{white} & \text{white} & \text{white} \\ \hline \text{white} & \text{shaded} & \text{white} & \text{white} \\ \hline \text{white} & \text{white} & \text{shaded} & \text{white} \\ \hline \text{white} & \text{white} & \text{white} & \text{shaded} \\ \hline \end{array} \end{array}$$

The idea is that by excising the diagonal blocks from \mathbf{A} , we obtain a rank-deficient matrix $\mathbf{A} - \underline{\mathbf{D}}$ that can be factored with block diagonal flanking matrices: $\mathbf{A} - \underline{\mathbf{D}} = \underline{\mathbf{U}} \tilde{\mathbf{A}} \underline{\mathbf{V}}^*$.

The inverse of a block-separable matrix can rapidly be constructed using the following simple variation of the classical Sherman-Morrison-Woodbury formula:

Lemma 3.1. *Suppose that \mathbf{A} is an $N \times N$ invertible matrix. Suppose further that K is a positive integer such that $K < N$, that \mathbf{A} admits the decomposition*

$$(3.5) \quad \begin{array}{ccccc} \mathbf{A} & = & \mathbf{U} & \tilde{\mathbf{A}} & \mathbf{V}^* \\ N \times N & & N \times K & K \times K & K \times N \end{array} + \begin{array}{c} \mathbf{D} \\ N \times N \end{array},$$

and that the matrices \mathbf{D} , $(\mathbf{V}^* \mathbf{D}^{-1} \mathbf{U})$, and $(\tilde{\mathbf{A}} + (\mathbf{V}^* \mathbf{D}^{-1} \mathbf{U})^{-1})$ are invertible. Then

$$(3.6) \quad \mathbf{A}^{-1} = \mathbf{E} (\tilde{\mathbf{A}} + \hat{\mathbf{D}})^{-1} \mathbf{F}^* + \mathbf{G},$$

where

$$(3.7) \quad \hat{\mathbf{D}} = (\mathbf{V}^* \mathbf{D}^{-1} \mathbf{U})^{-1},$$

$$(3.8) \quad \mathbf{E} = \mathbf{D}^{-1} \mathbf{U} \hat{\mathbf{D}},$$

$$(3.9) \quad \mathbf{F} = (\hat{\mathbf{D}} \mathbf{V}^* \mathbf{D}^{-1})^*,$$

$$(3.10) \quad \mathbf{G} = \mathbf{D}^{-1} - \mathbf{D}^{-1} \mathbf{U} \hat{\mathbf{D}} \mathbf{V}^* \mathbf{D}^{-1}.$$

When \mathbf{A} is block-separable, (3.5) holds with block diagonal matrices \mathbf{U} , \mathbf{V} , and \mathbf{D} . The matrices $\hat{\mathbf{D}}$, \mathbf{E} , \mathbf{F} , and \mathbf{G} can then be evaluated rapidly, and Lemma 3.1 can be said to reduce the task of inverting the $np \times np$ matrix \mathbf{A} , to the task of inverting the $kp \times kp$ matrix $\tilde{\mathbf{A}} + \hat{\mathbf{D}}$.

Proof of Lemma 3.1: Consider the equation

$$(3.11) \quad (\mathbf{U} \tilde{\mathbf{A}} \mathbf{V}^* + \mathbf{D}) \mathbf{q} = \mathbf{u}.$$

We will prove that (3.6) holds by proving that the solution \mathbf{q} of (3.11) is the right hand side of (3.6) applied to \mathbf{u} . First we set

$$(3.12) \quad \hat{\mathbf{q}} = \mathbf{V}^* \mathbf{q}.$$

Then (3.11) can be written

$$(3.13) \quad \mathbf{U} \tilde{\mathbf{A}} \hat{\mathbf{q}} + \mathbf{D} \mathbf{q} = \mathbf{u}.$$

Solving (3.13) for \mathbf{q} and inserting the result in (3.12), we obtain

$$(3.14) \quad (I + \underbrace{\mathbf{V}^* \mathbf{D}^{-1} \mathbf{U} \tilde{\mathbf{A}}}_{=\hat{\mathbf{D}}^{-1}}) \hat{\mathbf{q}} = \mathbf{V}^* \mathbf{D}^{-1} \mathbf{u}.$$

Multiplying (3.14) by $\hat{\mathbf{D}}$ we find that

$$(3.15) \quad (\hat{\mathbf{D}} + \tilde{\mathbf{A}}) \hat{\mathbf{q}} = \underbrace{\hat{\mathbf{D}} \mathbf{V}^* \mathbf{D}^{-1}}_{=\mathbf{F}^*} \mathbf{u}.$$

Now note that from (3.13) it also follows that

$$(3.16) \quad \mathbf{q} = -\mathbf{D}^{-1} \mathbf{U} \tilde{\mathbf{A}} \hat{\mathbf{q}} + \mathbf{D}^{-1} \mathbf{u}.$$

From (3.15) we know that

$$(3.17) \quad \tilde{\mathbf{A}} \hat{\mathbf{q}} = -\hat{\mathbf{D}} \hat{\mathbf{q}} + \mathbf{F}^* \mathbf{u}.$$

Inserting (3.17) into (3.16), we obtain

$$(3.18) \quad \mathbf{q} = -\mathbf{D}^{-1} \mathbf{U} (-\hat{\mathbf{D}} \hat{\mathbf{q}} + \mathbf{F}^* \mathbf{u}) + \mathbf{D}^{-1} \mathbf{u} = \underbrace{\mathbf{D}^{-1} \mathbf{U} \hat{\mathbf{D}}}_{=\mathbf{E}} \hat{\mathbf{q}} + \underbrace{(\mathbf{D}^{-1} - \mathbf{D}^{-1} \mathbf{U} \mathbf{F}^*)}_{=\mathbf{G}} \mathbf{u}.$$

Solving (3.15) for $\hat{\mathbf{q}}$ and inserting the result in (3.18), we obtain the expression (3.6). \square

The technique provided by the lemma is a useful tool in its own right. It can often reduce the cost of inverting an $N \times N$ matrix from the $O(N^3)$ cost of Gaussian elimination, to $O(N^{1.8})$, see Remark 3.1. Even more significant is that for many matrices, including the ones under consideration in this paper, the process described can be continued recursively which leads to an $O(N)$ inversion scheme. The required hierarchical representations of matrices are described in Section 4, and the $O(N)$ inversion scheme is given in Section 5.

Remark 3.1. To assess the computational cost of applying Lemma 3.1, suppose that \mathbf{A} is a BS matrix whose $p \times p$ blocks of size $n \times n$ satisfy (3.2). Then evaluating the factors $\hat{\mathbf{D}}$, \mathbf{E} , \mathbf{F} , and \mathbf{G} requires $O(p n^3)$ operations. Evaluating $(\tilde{\mathbf{A}} + \hat{\mathbf{D}})^{-1}$ requires $O(p^3 k^3)$ operations. The total cost T_{BS} of evaluating the factors in formula (3.6) therefore satisfies:

$$(3.19) \quad T_{\text{BS}} \sim p n^3 + p^3 k^3.$$

The cost (3.19) should be compared to the $O(p^3 n^3)$ cost of evaluating \mathbf{A}^{-1} directly. To elucidate the comparison, suppose temporarily that we are given a matrix \mathbf{A} of fixed size $N \times N$, and can choose how many blocks p we wish to partition it into. Then $n \approx N/p$, and the total cost satisfies

$$(3.20) \quad T_{\text{BS}} \sim p^{-2} N^3 + p^3 k^3.$$

If the rank of interaction k is independent of the block size, we can set $p \sim N^{3/5}$, whence

$$(3.21) \quad T_{\text{BS}} \sim k^3 N^{9/5}.$$

In practice, the numerical rank of the off-diagonal blocks typically increases slightly as the block size n is increased, but the increase tends to be moderate. For instance, for the matrices under consideration in this paper, one typically sees $k \sim \log(n) \sim \log(N/p)$. In such an environment, setting $p \sim N^{3/5} (\log N)^{-3/5}$ transforms (3.20) to, at worst,

$$T_{\text{BS}} \sim (\log N)^{6/5} N^{9/5}.$$

The calculations in this remark do not include the cost of actually constructing the factors \mathbf{U} , \mathbf{V} , \mathbf{D} . For a general matrix, this cost is $O(k N^2)$, but for the matrices under consideration in this paper, techniques with better asymptotic complexity are described in Section 6.

We close by showing that when the matrix \mathbf{A} is symmetric positive definite, the assumption in Lemma 3.1 that certain intermediate matrices are invertible can be omitted:

Corollary 3.2. *Let \mathbf{A} be a symmetric positive definite (spd) matrix that admits a factorization*

$$(3.22) \quad \begin{array}{ccccc} \mathbf{A} & = & \mathbf{U} & \tilde{\mathbf{A}} & \mathbf{U}^* & + & \mathbf{D}, \\ N \times N & & N \times K & K \times K & K \times N & & N \times N \end{array}$$

where $\ker(\mathbf{U}) = \{\mathbf{0}\}$ and \mathbf{D} is a block diagonal submatrix of \mathbf{A} . Then the matrices \mathbf{D} , $(\mathbf{U}^* \mathbf{D}^{-1} \mathbf{U})$, and $\tilde{\mathbf{A}} + (\mathbf{U}^* \mathbf{D}^{-1} \mathbf{U})^{-1}$ are spd (and hence invertible).

Proof of Corollary 3.2: That \mathbf{D} is spd follows immediately from the fact that it is a block diagonal submatrix of a spd matrix.

To show that $\mathbf{U}^* \mathbf{D}^{-1} \mathbf{U}$ is spd we pick any $\mathbf{x} \neq \mathbf{0}$, set $\mathbf{y} = \mathbf{D}^{-1} \mathbf{U} \mathbf{x}$, observe that $\mathbf{y} \neq \mathbf{0}$ since $\ker(\mathbf{U}) = \{\mathbf{0}\}$, and then we find that $\langle \mathbf{U}^* \mathbf{D}^{-1} \mathbf{U} \mathbf{x}, \mathbf{x} \rangle = \langle \mathbf{D}^{-1} \mathbf{U}, \mathbf{U} \mathbf{x} \rangle = \langle \mathbf{y}, \mathbf{D} \mathbf{y} \rangle > 0$ since \mathbf{D} is spd.

It remains only to prove that $\tilde{\mathbf{A}} + (\mathbf{U}^* \mathbf{D}^{-1} \mathbf{U})^{-1}$ is spd. To this end, define $\hat{\mathbf{D}}$ and \mathbf{E} via

$$\begin{aligned} \hat{\mathbf{D}} &= (\mathbf{U}^* \mathbf{D}^{-1} \mathbf{U})^{-1} \\ \mathbf{E} &= \mathbf{D}^{-1} \mathbf{U} \hat{\mathbf{D}}. \end{aligned}$$

Then

$$\begin{aligned} \tilde{\mathbf{A}} + \hat{\mathbf{D}} &= \hat{\mathbf{D}} \left(\hat{\mathbf{D}}^{-1} \tilde{\mathbf{A}} \hat{\mathbf{D}}^{-1} + \hat{\mathbf{D}}^{-1} \right) \hat{\mathbf{D}} = \hat{\mathbf{D}} \left(\mathbf{U}^* \mathbf{D}^{-1} \mathbf{U} \tilde{\mathbf{A}} \mathbf{U}^* \mathbf{D}^{-1} \mathbf{U} + \hat{\mathbf{D}}^{-1} \right) \hat{\mathbf{D}} \\ &= \hat{\mathbf{D}} \left(\mathbf{U}^* \mathbf{D}^{-1} (\mathbf{A} - \mathbf{D}) \mathbf{D}^{-1} \mathbf{U} + \mathbf{U}^* \mathbf{D}^{-1} \mathbf{U} \right) \hat{\mathbf{D}} = \hat{\mathbf{D}} \mathbf{U}^* \mathbf{D}^{-1} \mathbf{A} \mathbf{D}^{-1} \mathbf{U} \hat{\mathbf{D}} = \mathbf{E}^* \mathbf{A} \mathbf{E}. \end{aligned}$$

That $\tilde{\mathbf{A}} + (\mathbf{U}^* \mathbf{D}^{-1} \mathbf{U})^{-1}$ is spd now follows since $\ker(\mathbf{E}) = \{\mathbf{0}\}$ and \mathbf{A} is spd. \square

Remark 3.2. The proof of Corollary 3.2 demonstrates that the stability of the method can readily be assessed by tracking the conditioning of the matrices \mathbf{E} . If these matrices are close to orthogonal (i.e. all their singular values are similar in magnitude) then the compressed matrix $\tilde{\mathbf{A}} + \hat{\mathbf{D}}$ has about the same distribution of singular values as \mathbf{A} since $\tilde{\mathbf{A}} + \hat{\mathbf{D}} = \mathbf{E}^* \mathbf{A} \mathbf{E}$.

4. HIERARCHICALLY BLOCK SEPARABLE MATRICES

In this section, we define what it means for an $N \times N$ matrix \mathbf{A} to be HBS. Section 4.1 informally describes the basic ideas. Section 4.2 describes a simple binary tree of subsets of the index vector $[1, 2, \dots, N]$. Section 4.3 provides a formal definition of the HBS property. Section 4.4 describes how an HBS matrix can be expressed a telescoping factorization. Section 4.5 describes an $O(N)$ procedure for applying an HBS matrix to a vector.

4.1. Heuristic description. The HBS property first of all requires \mathbf{A} to be BS. Supposing that \mathbf{A} consists of 8×8 blocks, this means that \mathbf{A} admits a factorization, *cf.* (3.4):

$$(4.1) \quad \mathbf{A} = \underline{\mathbf{U}}^{(3)} \tilde{\mathbf{A}}^{(3)} (\underline{\mathbf{V}}^{(3)})^* + \underline{\mathbf{D}}^{(3)}$$

The superscripts on the right-hand side of (4.1) indicate that the factorization is at “level 3.” We next require the smaller matrix $\tilde{\mathbf{A}}^{(3)}$ to be BS, and to admit the analogous factorization:

$$(4.2) \quad \tilde{\mathbf{A}}^{(3)} = \underline{\mathbf{U}}^{(2)} \tilde{\mathbf{A}}^{(2)} (\underline{\mathbf{V}}^{(2)})^* + \underline{\mathbf{B}}^{(2)}$$

In forming (4.2), we reblocked the matrix $\tilde{\mathbf{A}}^{(3)}$ by merging blocks in groups of four. The purpose is to “reintroduce” rank deficiencies in the off-diagonal blocks. In the final step in the hierarchy, we require that upon reblocking, $\tilde{\mathbf{A}}^{(2)}$ is BS and admits a factorization:

$$(4.3) \quad \tilde{\mathbf{A}}^{(2)} = \underline{\mathbf{U}}^{(1)} \tilde{\mathbf{A}}^{(1)} (\underline{\mathbf{V}}^{(1)})^* + \underline{\mathbf{B}}^{(1)}$$

Combining (4.1), (4.2), and (4.3), we find that \mathbf{A} can be expressed as

$$(4.4) \quad \mathbf{A} = \underline{\mathbf{U}}^{(3)} (\underline{\mathbf{U}}^{(2)} (\underline{\mathbf{U}}^{(1)} \underline{\mathbf{B}}^{(0)} (\underline{\mathbf{V}}^{(1)})^* + \underline{\mathbf{B}}^{(1)}) (\underline{\mathbf{V}}^{(2)})^* + \underline{\mathbf{B}}^{(2)}) (\underline{\mathbf{V}}^{(3)})^* + \underline{\mathbf{D}}^{(3)}.$$

The block structure of the right hand side of (4.4) is:

In other words, the HBS property lets us completely represent a matrix in terms of certain block diagonal factors. The total cost of storing these factors is $O(Nk)$, and the format is an example of a so-called “data-sparse” representation of the matrix.

Remark 4.1. In describing the HBS property, we assumed that all off-diagonal blocks at all levels have the same rank k . We do this for notational clarity only. In practice, the minimal rank tends to vary slightly from block to block, and to moderately increase on the coarser levels. In the numerical examples in Section 7, all codes determine the rank adaptively for each block.

4.2. Tree structure. The HBS representation of an $N \times N$ matrix \mathbf{A} is based on a partition of the index vector $I = [1, 2, \dots, N]$ into a binary tree structure. For simplicity, we limit attention to binary tree structures in which every level is fully populated. We let I form the root of the tree, and give it the index 1, $I_1 = I$. We next split the root into two roughly equi-sized vectors I_2 and I_3 so that $I_1 = I_2 \cup I_3$. The full tree is then formed by continuing to subdivide any interval that holds more than some preset fixed number n of indices. We use the integers $\ell = 0, 1, \dots, L$ to label the different levels, with 0 denoting the coarsest level. A *leaf* is a node corresponding

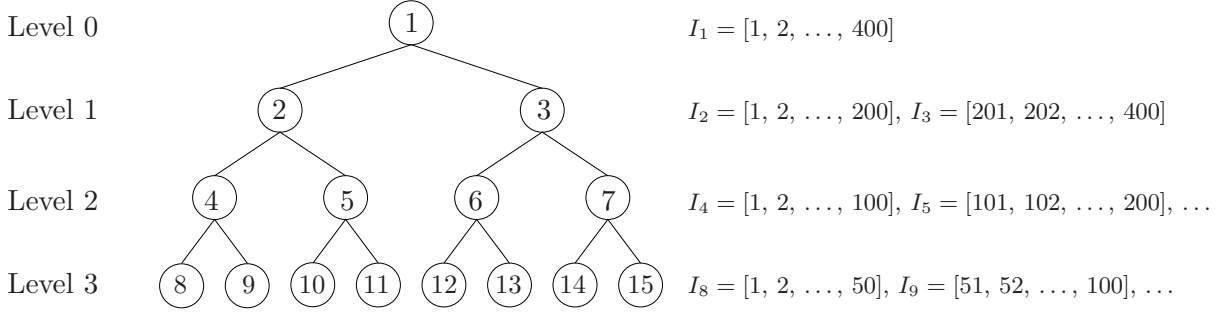


FIGURE 4.1. Numbering of nodes in a fully populated binary tree with $L = 3$ levels. The root is the original index vector $I = I_1 = [1, 2, \dots, 400]$.

to a vector that never got split. For a non-leaf node τ , its *children* are the two boxes σ_1 and σ_2 such that $I_\tau = I_{\sigma_1} \cup I_{\sigma_2}$, and τ is then the *parent* of σ_1 and σ_2 . Two boxes with the same parent are called *siblings*. These definitions are illustrated in Figure 4.1

Remark 4.2. The HBS format works with a broad range of different tree structures. It is permissible to split a node into more than two children if desirable, to distribute the points in an index set unevenly among its children, to split only some nodes on a given level, etc. This flexibility is essential when A approximates a non-uniformly discretized integral operator; in this case, the partition tree it constructed based on a geometric subdivision of the domain of integration. The *spatial* geometry of a box then dictates whether and how it is to be split. The algorithms of this paper can easily be modified to accommodate general trees.

4.3. Definition of the HBS property. We are now prepared to rigorously define what it means for an $N \times N$ matrix A to be *hierarchically block seperable* with respect to a given binary tree \mathcal{T} that partitions the index vector $J = [1, 2, \dots, N]$. For simplicity, we suppose that the tree has L fully populated levels, and that for every leaf node τ , the index vector I_τ holds precisely n points, so that $N = n 2^L$. Then A is HBS with block rank k if the following two conditions hold:

(1) *Assumption on ranks of off-diagonal blocks at the finest level:* For any two distinct leaf nodes τ and τ' , define the $n \times n$ matrix

$$(4.5) \quad A_{\tau, \tau'} = A(I_\tau, I_{\tau'}).$$

Then there must exist matrices U_τ , $V_{\tau'}$, and $\tilde{A}_{\tau, \tau'}$ such that

$$(4.6) \quad \begin{array}{ccccc} A_{\tau, \tau'} & = & U_\tau & \tilde{A}_{\tau, \tau'} & V_{\tau'}^* \\ n \times n & & n \times k & k \times k & k \times n \end{array}$$

(2) *Assumption on ranks of off-diagonal blocks on level $\ell = L - 1, L - 2, \dots, 1$:* The rank assumption at level ℓ is defined in terms of the blocks constructed on the next finer level $\ell + 1$: For any distinct nodes τ and τ' on level ℓ with children σ_1, σ_2 and σ'_1, σ'_2 , respectively, define

$$(4.7) \quad A_{\tau, \tau'} = \begin{bmatrix} \tilde{A}_{\sigma_1, \sigma'_1} & \tilde{A}_{\sigma_1, \sigma'_2} \\ \tilde{A}_{\sigma_2, \sigma'_1} & \tilde{A}_{\sigma_2, \sigma'_2} \end{bmatrix}.$$

Then there must exist matrices U_τ , $V_{\tau'}$, and $\tilde{A}_{\tau, \tau'}$ such that

$$(4.8) \quad \begin{array}{ccccc} A_{\tau, \tau'} & = & U_\tau & \tilde{A}_{\tau, \tau'} & V_{\tau'}^* \\ 2k \times 2k & & 2k \times k & k \times k & k \times 2k \end{array}$$

	Name:	Size:	Function:
For each leaf node τ :	\mathbf{D}_τ	$n \times n$	The diagonal block $\mathbf{A}(I_\tau, I_\tau)$.
	\mathbf{U}_τ	$n \times k$	Basis for the columns in the blocks in row τ .
	\mathbf{V}_τ	$n \times k$	Basis for the rows in the blocks in column τ .
For each parent node τ :	\mathbf{B}_τ	$2k \times 2k$	Interactions between the children of τ .
	\mathbf{U}_τ	$2k \times k$	Basis for the columns in the (reduced) blocks in row τ .
	\mathbf{V}_τ	$2k \times k$	Basis for the rows in the (reduced) blocks in column τ .

FIGURE 4.2. An HBS matrix \mathbf{A} associated with a tree \mathcal{T} is fully specified if the factors listed above are provided.

The two points above complete the definition. An HBS matrix is now fully described if the basis matrices \mathbf{U}_τ and \mathbf{V}_τ are provided for each node τ , and in addition, we are for each leaf τ given the $n \times n$ matrix

$$(4.9) \quad \mathbf{D}_\tau = \mathbf{A}(I_\tau, I_\tau),$$

and for each parent node τ with children σ_1 and σ_2 we are given the $2k \times 2k$ matrix

$$(4.10) \quad \mathbf{B}_\tau = \begin{bmatrix} 0 & \tilde{\mathbf{A}}_{\sigma_1, \sigma_2} \\ \tilde{\mathbf{A}}_{\sigma_2, \sigma_1} & 0 \end{bmatrix}.$$

Observe in particular that the matrices $\tilde{\mathbf{A}}_{\sigma_1, \sigma_2}$ are only required when $\{\sigma_1, \sigma_2\}$ forms a sibling pair. Figure 4.2 summarizes the required matrices.

Remark 4.3. The definition of the HBS property given in this section is flexible in the sense that we do not enforce any conditions on the factors \mathbf{U}_τ , \mathbf{V}_τ , and $\tilde{\mathbf{A}}_{\tau, \tau'}$ other than that (4.6) and (4.8) must hold. For purposes of numerical stability, further conditions are sometimes imposed. The perhaps strongest such condition is to require the matrices \mathbf{U}_τ and $\mathbf{V}_{\tau'}$ in (4.6) and (4.8) be orthonormal, see e.g. [26, 6] (one can in this case require that the matrices $\tilde{\mathbf{A}}_{\tau, \tau'}$ be diagonal, so that (4.6) and (4.8) become singular value decompositions). If a “general” HBS matrix is given, it can easily be converted to this more restrictive format via, e.g., Algorithm 1. A choice that we have found highly convenient is to require (4.6) and (4.8) to be interpolatory decompositions (see Section 2.2). Then every \mathbf{U}_τ and $\mathbf{V}_{\tau'}$ contains a $k \times k$ identity matrix (which greatly accelerates computations), and each $\tilde{\mathbf{A}}_{\tau, \tau'}$ is a submatrix of the original matrix \mathbf{A} .

Remark 4.4. The definition (4.6) transparently describes the functions of the basis matrices \mathbf{U}_τ and \mathbf{V}_τ whenever τ is a leaf node. The definition (4.8) of basis matrices for non-leaf nodes is perhaps less intuitive. Their meaning can be made clearer by defining what we call “extended” basis matrices $\mathbf{U}_\tau^{\text{extend}}$ and $\mathbf{V}_\tau^{\text{extend}}$. For a leaf node τ we simply set

$$\mathbf{U}_\tau^{\text{extend}} = \mathbf{U}_\tau, \quad \text{and} \quad \mathbf{V}_\tau^{\text{extend}} = \mathbf{V}_\tau.$$

For a parent node τ with children σ_1 and σ_2 , we set

$$\mathbf{U}_\tau^{\text{extend}} = \begin{bmatrix} \mathbf{U}_{\sigma_1}^{\text{extend}} & 0 \\ 0 & \mathbf{U}_{\sigma_2}^{\text{extend}} \end{bmatrix} \mathbf{U}_\tau, \quad \text{and} \quad \mathbf{V}_\tau^{\text{extend}} = \begin{bmatrix} \mathbf{V}_{\sigma_1}^{\text{extend}} & 0 \\ 0 & \mathbf{V}_{\sigma_2}^{\text{extend}} \end{bmatrix} \mathbf{V}_\tau.$$

Then for any distinct nodes τ and τ' on level ℓ , we have

$$\mathbf{A}(I_\tau, I_{\tau'})_{n 2^{L-\ell} \times n 2^{L-\ell}} = \mathbf{U}_\tau^{\text{extend}}_{n 2^{L-\ell} \times k} \tilde{\mathbf{A}}_{\tau, \tau'}_{k \times k} (\mathbf{V}_{\tau'}^{\text{extend}})^*_{k \times n 2^{L-\ell}}.$$

ALGORITHM 1 (reformatting an HBS matrix)

Given the factors $\mathbf{U}_\tau, \mathbf{V}_\tau, \tilde{\mathbf{A}}_{\sigma_1, \sigma_2}, \mathbf{D}_\tau$ of an HBS matrix in general format, this algorithm computes new factors (that overwrite the old) such that all \mathbf{U}_τ and \mathbf{V}_τ are orthonormal, and all $\tilde{\mathbf{A}}_{\sigma_1, \sigma_2}$ are diagonal.

loop over levels, finer to coarser, $\ell = L - 1, L - 2, \dots, 0$

loop over all parent boxes τ on level ℓ ,

 Let σ_1 and σ_2 denote the children of τ .

$[\mathbf{U}_1, \mathbf{R}_1] = \mathbf{qr}(\mathbf{U}_{\sigma_1})$.

$[\mathbf{V}_1, \mathbf{S}_1] = \mathbf{qr}(\mathbf{V}_{\sigma_1})$.

$[\mathbf{X}_1, \Sigma_{12}, \mathbf{Y}_2] = \mathbf{svd}(\mathbf{R}_1 \tilde{\mathbf{A}}_{\sigma_1, \sigma_2} \mathbf{S}_2^*)$.

$\mathbf{U}_{\sigma_1} \leftarrow \mathbf{U}_1 \mathbf{X}_1$.

$\mathbf{V}_{\sigma_1} \leftarrow \mathbf{V}_1 \mathbf{Y}_1$.

$\mathbf{B}_{\sigma_1 \sigma_2} \leftarrow \Sigma_{12}$.

if $\ell > 0$

$$\mathbf{U}_\tau \leftarrow \begin{bmatrix} \mathbf{X}_1^* \mathbf{R}_1 & 0 \\ 0 & \mathbf{X}_2^* \mathbf{R}_2 \end{bmatrix} \mathbf{U}_\tau.$$

end if

end loop

end loop

$[\mathbf{U}_2, \mathbf{R}_2] = \mathbf{qr}(\mathbf{U}_{\sigma_2})$.

$[\mathbf{V}_2, \mathbf{S}_2] = \mathbf{qr}(\mathbf{V}_{\sigma_2})$.

$[\mathbf{X}_2, \Sigma_{21}, \mathbf{Y}_1] = \mathbf{svd}(\mathbf{R}_2 \tilde{\mathbf{A}}_{\sigma_2, \sigma_1} \mathbf{S}_1^*)$.

$\mathbf{U}_{\sigma_2} \leftarrow \mathbf{U}_2 \mathbf{X}_2$.

$\mathbf{V}_{\sigma_2} \leftarrow \mathbf{V}_2 \mathbf{Y}_1$.

$\mathbf{B}_{\sigma_2 \sigma_1} \leftarrow \Sigma_{21}$.

$$\mathbf{V}_\tau \leftarrow \begin{bmatrix} \mathbf{Y}_1^* \mathbf{S}_1 & 0 \\ 0 & \mathbf{Y}_2^* \mathbf{S}_2 \end{bmatrix} \mathbf{V}_\tau.$$

4.4. Telescoping factorizations. In the heuristic description of the HBS property in Section 4.1, we claimed that any HBS matrix can be expressed as a telescoping factorization with block diagonal factors. We will now formalize this claim. Given the matrices defined in Section 4.3 (and summarized in Figure 4.2), we define the following block diagonal factors:

$$(4.11) \quad \underline{\mathbf{D}}^{(\ell)} = \text{diag}(\mathbf{D}_\tau : \tau \text{ is a box on level } \ell), \quad \ell = 0, 1, \dots, L,$$

$$(4.12) \quad \underline{\mathbf{U}}^{(\ell)} = \text{diag}(\mathbf{U}_\tau : \tau \text{ is a box on level } \ell), \quad \ell = 1, 2, \dots, L,$$

$$(4.13) \quad \underline{\mathbf{V}}^{(\ell)} = \text{diag}(\mathbf{V}_\tau : \tau \text{ is a box on level } \ell), \quad \ell = 1, 2, \dots, L,$$

$$(4.14) \quad \underline{\mathbf{B}}^{(\ell)} = \text{diag}(\mathbf{B}_\tau : \tau \text{ is a box on level } \ell), \quad \ell = 0, 1, \dots, L - 1, .$$

Furthermore, we let $\tilde{\mathbf{A}}^{(\ell)}$ denote the block matrix whose diagonal blocks are zero, and whose off-diagonal blocks are the blocks $\tilde{\mathbf{A}}_{\tau, \tau'}$ for all distinct τ, τ' on level ℓ . With these definitions,

$$(4.15) \quad \begin{matrix} \mathbf{A} \\ n 2^L \times n 2^L \end{matrix} = \begin{matrix} \underline{\mathbf{U}}^{(L)} \\ n 2^L \times k 2^L \end{matrix} \begin{matrix} \tilde{\mathbf{A}}^{(L)} \\ k 2^L \times k 2^L \end{matrix} \begin{matrix} (\underline{\mathbf{V}}^{(L)})^* \\ k 2^L \times n 2^L \end{matrix} + \begin{matrix} \underline{\mathbf{D}}^{(L)} \\ n 2^L \times n 2^L \end{matrix};$$

for $\ell = L - 1, L - 2, \dots, 1$ we have

$$(4.16) \quad \begin{matrix} \tilde{\mathbf{A}}^{(\ell+1)} \\ k 2^{\ell+1} \times k 2^{\ell+1} \end{matrix} = \begin{matrix} \underline{\mathbf{U}}^{(\ell)} \\ k 2^{\ell+1} \times k 2^\ell \end{matrix} \begin{matrix} \tilde{\mathbf{A}}^{(\ell)} \\ k 2^\ell \times k 2^\ell \end{matrix} \begin{matrix} (\underline{\mathbf{V}}^{(\ell)})^* \\ k 2^\ell \times k 2^{\ell+1} \end{matrix} + \begin{matrix} \underline{\mathbf{B}}^{(\ell)} \\ k 2^{\ell+1} \times k 2^{\ell+1} \end{matrix};$$

and finally

$$(4.17) \quad \tilde{\mathbf{A}}^{(1)} = \underline{\mathbf{B}}^{(0)}.$$

4.5. Matrix-vector multiplication. The telescoping factorizations in Section 4.4 easily translate into a formula for evaluating the matrix-vector product $\mathbf{u} = \mathbf{A} \mathbf{q}$ once all factors in an HBS representation have been provided. The resulting algorithm has computational complexity $O(Nk)$ (assuming that $n = O(k)$), and is given as Algorithm 2.

ALGORITHM 2 (HBS matrix-vector multiply)

Given a vector \mathbf{q} and a matrix \mathbf{A} in HBS format, compute $\mathbf{u} = \mathbf{A} \mathbf{q}$.

loop over all leaf boxes τ

$$\hat{\mathbf{q}}_\tau = \mathbf{V}_\tau^* \mathbf{q}(I_\tau).$$

end loop

loop over levels, finer to coarser, $\ell = L - 1, L - 2, \dots, 1$

loop over all parent boxes τ on level ℓ ,

Let σ_1 and σ_2 denote the children of τ .

$$\hat{\mathbf{q}}_\tau = \mathbf{V}_\tau^* \begin{bmatrix} \hat{\mathbf{q}}_{\sigma_1} \\ \hat{\mathbf{q}}_{\sigma_2} \end{bmatrix}.$$

end loop

end loop

$$\hat{\mathbf{u}}_1 = 0$$

loop over all levels, coarser to finer, $\ell = 1, 2, \dots, L - 1$

loop over all parent boxes τ on level ℓ

Let σ_1 and σ_2 denote the children of τ .

$$\begin{bmatrix} \hat{\mathbf{u}}_{\sigma_1} \\ \hat{\mathbf{u}}_{\sigma_2} \end{bmatrix} = \mathbf{U}_\tau \hat{\mathbf{u}}_\tau + \begin{bmatrix} 0 & \mathbf{B}_{\sigma_1, \sigma_2} \\ \mathbf{B}_{\sigma_2, \sigma_1} & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{q}}_{\sigma_1} \\ \hat{\mathbf{q}}_{\sigma_2} \end{bmatrix}.$$

end loop

end loop

loop over all leaf boxes τ

$$\mathbf{u}(I_\tau) = \mathbf{U}_\tau \hat{\mathbf{u}}_\tau + \mathbf{D}_\tau \mathbf{q}(I_\tau).$$

end loop

5. INVERSION OF HIERARCHICALLY BLOCK SEPARABLE MATRICES

In this section, we describe an algorithm for inverting an HBS matrix \mathbf{A} . The algorithm is exact (in the absence of round-off errors) and has asymptotic complexity $O(Nk^2)$. It is summarized as Algorithm 3, and as the description indicates, it is very simple to implement. The output of Algorithm 3 is a set of factors in a data-sparse representation of \mathbf{A}^{-1} which can be applied to a given vector via Algorithm 4. Technically, the scheme consists of recursive application of Lemma 3.1, and its derivation is given in the proof of the following theorem:

Theorem 5.1. *Let \mathbf{A} be an invertible $N \times N$ HBS matrix with block-rank k . Suppose that at the finest level of the tree structure, there are 2^L leaves that each holds n points (so that $N = n2^L$), and that $n \leq 2k$. Then a data-sparse representation of \mathbf{A}^{-1} that is exact up to rounding errors can be computed in $O(Nk^2)$ operations via a process given as Algorithm 3, provided that none of the matrices that need to be inverted is singular. The computed inverse can be applied to a vector in $O(Nk)$ operations via Algorithm 4.*

Proof: We can according to equation (4.15) express an HBS matrix as

$$\mathbf{A} = \underline{\mathbf{U}}^{(L)} \tilde{\mathbf{A}}^{(L)} (\underline{\mathbf{V}}^{(L)})^* + \underline{\mathbf{D}}^{(L)}.$$

Lemma 3.1 immediately applies, and we find that

$$(5.1) \quad \mathbf{A}^{-1} = \underline{\mathbf{E}}^{(L)} (\tilde{\mathbf{A}}^{(L)} + \hat{\underline{\mathbf{D}}}^{(L)})^{-1} (\underline{\mathbf{F}}^{(L)})^* + \underline{\mathbf{G}}^{(L)},$$

where $\underline{\mathbf{E}}^{(L)}$, $\underline{\mathbf{F}}^{(L)}$, $\hat{\underline{\mathbf{D}}}^{(L)}$ and $\underline{\mathbf{G}}^{(L)}$ are defined via (3.7), (3.8), (3.9), and (3.10).

To move to the next coarser level, we set $\ell = L - 1$ in formula (4.16) whence

$$(5.2) \quad \tilde{\mathbf{A}}^{(L)} + \hat{\underline{\mathbf{D}}}^{(L)} = \underline{\mathbf{U}}^{(L-1)} \tilde{\mathbf{A}}^{(L-1)} (\underline{\mathbf{V}}^{(L-1)})^* + \underline{\mathbf{B}}^{(L-1)} + \hat{\underline{\mathbf{D}}}^{(L)}.$$

We define

$$\tilde{\underline{\mathbf{D}}}^{(L-1)} = \underline{\mathbf{B}}^{(L-1)} + \hat{\underline{\mathbf{D}}}^{(L)} = \left[\begin{array}{cc|cc|cc|c} \hat{\mathbf{D}}_{\tau_1} & \mathbf{B}_{\tau_1 \tau_2} & 0 & 0 & 0 & 0 & \cdots \\ \mathbf{B}_{\tau_1 \tau_2} & \hat{\mathbf{D}}_{\tau_2} & 0 & 0 & 0 & 0 & \cdots \\ \hline 0 & 0 & \hat{\mathbf{D}}_{\tau_3} & \mathbf{B}_{\tau_3 \tau_4} & 0 & 0 & \cdots \\ 0 & 0 & \mathbf{B}_{\tau_3 \tau_4} & \hat{\mathbf{D}}_{\tau_4} & 0 & 0 & \cdots \\ \hline 0 & 0 & 0 & 0 & \hat{\mathbf{D}}_{\tau_5} & \mathbf{B}_{\tau_5 \tau_6} & \cdots \\ 0 & 0 & 0 & 0 & \mathbf{B}_{\tau_5 \tau_6} & \hat{\mathbf{D}}_{\tau_6} & \cdots \\ \hline \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \end{array} \right],$$

where $\{\tau_1, \tau_2, \dots, \tau_{2L}\}$ is a list of the boxes on level L . Equation (5.2) then takes the form

$$(5.3) \quad \tilde{\mathbf{A}}^{(L)} + \hat{\underline{\mathbf{D}}}^{(L)} = \underline{\mathbf{U}}^{(L-1)} \tilde{\mathbf{A}}^{(L-1)} (\underline{\mathbf{V}}^{(L-1)})^* + \tilde{\underline{\mathbf{D}}}^{(L-1)}.$$

(We note that in (5.3), the terms on the left hand side are block matrices with $2^L \times 2^L$ blocks, each of size $k \times k$, whereas the terms on the right hand side are block matrices with $2^{L-1} \times 2^{L-1}$ blocks, each of size $2k \times 2k$.) Now Lemma 3.1 applies to (5.3) and we find that

$$(\tilde{\mathbf{A}}^{(L)} + \hat{\underline{\mathbf{D}}}^{(L)})^{-1} = \underline{\mathbf{E}}^{(L-1)} (\tilde{\mathbf{A}}^{(L-1)} + \hat{\underline{\mathbf{D}}}^{(L-1)})^{-1} (\underline{\mathbf{F}}^{(L-1)})^* + \underline{\mathbf{G}}^{(L-1)},$$

where $\underline{\mathbf{E}}^{(L-1)}$, $\underline{\mathbf{F}}^{(L-1)}$, $\hat{\underline{\mathbf{D}}}^{(L-1)}$ and $\underline{\mathbf{G}}^{(L-1)}$ are defined via (3.7), (3.8), (3.9), and (3.10).

The process by which we went from step L to step $L - 1$ is then repeated to move up to coarser and coarser levels. With each step, the size of the matrix to be inverted is cut in half. Once we get to the top level, we are left with the task of inverting the matrix

$$(5.4) \quad \tilde{\mathbf{A}}^{(1)} + \hat{\underline{\mathbf{D}}}^{(1)} = \begin{bmatrix} \hat{\mathbf{D}}_2 & \mathbf{B}_{2,3} \\ \mathbf{B}_{3,2} & \hat{\mathbf{D}}_3 \end{bmatrix}$$

The matrix in (5.4) is of size $2k \times 2k$, and we use brute force to evaluate

$$\mathbf{G}^{(0)} = \mathbf{G}_1 = \begin{bmatrix} \hat{\mathbf{D}}_2 & \mathbf{B}_{2,3} \\ \mathbf{B}_{3,2} & \hat{\mathbf{D}}_3 \end{bmatrix}^{-1}.$$

To calculate the cost of the inversion scheme described, we note that in order to compute the matrices $\underline{\mathbf{E}}^{(\ell)}$, $\underline{\mathbf{F}}^{(\ell)}$, $\underline{\mathbf{G}}^{(\ell)}$, and $\hat{\underline{\mathbf{D}}}^{(\ell)}$ on level ℓ , we need to perform dense matrix operations on 2^ℓ blocks, each of size at most $2k \times 2k$. Since the leaf boxes each hold at most $2k$ points, so that $2^{L+1}k \leq N$, the total cost is

$$\text{COST} \sim \sum_{\ell=1}^L 2^\ell 8k^3 \sim 2^{L+4} k^3 \sim N k^2.$$

This completes the proof. \square

Remark 5.1. Algorithm 3 produces a representation of \mathbf{A}^{-1} that is not exactly in HBS form since the matrices \mathbf{G}_τ do not have zero diagonal blocks like the matrices \mathbf{B}_τ , cf. (4.10). However, a simple technique given as Algorithm 5 converts the factorization provided by Algorithm 3 into a standard HBS factorization. If a factorization in which the expansion matrices are all orthonormal is sought, then further post-processing via Algorithm 1 will do the job.

ALGORITHM 3 (inversion of an HBS matrix)

```

loop over all levels, finer to coarser,  $\ell = L, L - 1, \dots, 1$ 
  loop over all boxes  $\tau$  on level  $\ell$ ,
    if  $\tau$  is a leaf node
       $\tilde{D}_\tau = D_\tau$ 
    else
      Let  $\sigma_1$  and  $\sigma_2$  denote the children of  $\tau$ .
       $\tilde{D}_\tau = \begin{bmatrix} \hat{D}_{\sigma_1} & B_{\sigma_1, \sigma_2} \\ B_{\sigma_2, \sigma_1} & \hat{D}_{\sigma_2} \end{bmatrix}$ 
    end if
     $\hat{D}_\tau = (V_\tau^* \tilde{D}_\tau^{-1} U_\tau)^{-1}$ .
     $E_\tau = \tilde{D}_\tau^{-1} U_\tau \hat{D}_\tau$ .
     $F_\tau^* = \hat{D}_\tau V_\tau^* \tilde{D}_\tau^{-1}$ .
     $G_\tau = \hat{D}_\tau - \tilde{D}_\tau^{-1} U_\tau \hat{D}_\tau V_\tau^* \tilde{D}_\tau^{-1}$ .
  end loop
end loop
 $G_1 = \begin{bmatrix} \hat{D}_2 & B_{2,3} \\ B_{3,2} & \hat{D}_3 \end{bmatrix}^{-1}$ .

```

Remark 5.2. Algorithm 3 provides four formulas for the matrices $\{E_\tau, F_\tau, G_\tau, \hat{D}_\tau\}_\tau$. The task of actually computing the matrices can be accelerated in two ways: (1) Several of the matrix-matrix products in the formulas are recurring, and the computation should be organized so that each matrix-matrix product is evaluated only once. (2) When interpolatory decompositions are used, multiplications involving the matrices U_τ and V_τ can be accelerated by exploiting that they each contain a $k \times k$ identity matrix.

Remark 5.3. The assumption in Theorem 5.1 that none of the matrices to be inverted is singular is undesirable. When A is spd, this assumption can be done away with (cf. Corollary 3.2), and it can further be proved that the inversion process is numerically stable. When A is non-symmetric, the intermediate matrices often become ill-conditioned. We have *empirically* observed that if we enforce that $U_\tau = V_\tau$ for every node (the procedure for doing this for non-symmetric matrices is described in Remark 6.1), then the method is remarkably stable, but we do not have any supporting theory.

Remark 5.4. The assumption in Theorem 5.1 that the block-rank k remains constant across all levels is slightly unnatural. In applications to integral equations on 1D domain, one often finds that the rank of interaction depends logarithmically on the number of points in a block. It is shown in [21] that inclusion of such logarithmic growth of the interaction ranks does not change the $O(N)$ total complexity.

6. COMPUTING THE HBS REPRESENTATION OF A BOUNDARY INTEGRAL OPERATOR

Section 4 describes a particular way of representing a class of “compressible” matrices in a hierarchical structure of block-diagonal matrices. For any matrix whose HBS rank is k , these factors can via straight-forward means be computed in $O(N^2 k)$ operations. In this section, we describe an $O(N k^2)$ technique for computing an HBS representation of the matrix resulting upon Nyström discretization of a BIE.

ALGORITHM 4 (application of inverse)

Given a vector \mathbf{u} , compute $\mathbf{q} = \mathbf{A}^{-1} \mathbf{u}$ using the compressed representation of \mathbf{A}^{-1} resulting from Algorithm 3.

loop over all leaf boxes τ

$$\hat{\mathbf{u}}_\tau = \mathbf{F}_\tau^* \mathbf{u}(I_\tau).$$

end loop

loop over all levels, finer to coarser, $\ell = L, L-1, \dots, 1$

loop over all parent boxes τ on level ℓ ,

 Let σ_1 and σ_2 denote the children of τ .

$$\hat{\mathbf{u}}_\tau = \mathbf{F}_\tau^* \begin{bmatrix} \hat{\mathbf{u}}_{\sigma_1} \\ \hat{\mathbf{u}}_{\sigma_2} \end{bmatrix}.$$

end loop

end loop

$$\begin{bmatrix} \hat{\mathbf{q}}_2 \\ \hat{\mathbf{q}}_3 \end{bmatrix} = \hat{\mathbf{G}}_1 \begin{bmatrix} \hat{\mathbf{u}}_2 \\ \hat{\mathbf{u}}_3 \end{bmatrix}.$$

loop over all levels, coarser to finer, $\ell = 1, 2, \dots, L-1$

loop over all parent boxes τ on level ℓ

 Let σ_1 and σ_2 denote the children of τ .

$$\begin{bmatrix} \hat{\mathbf{q}}_{\sigma_1} \\ \hat{\mathbf{q}}_{\sigma_2} \end{bmatrix} = \mathbf{E}_\tau \hat{\mathbf{u}}_\tau + \mathbf{G}_\tau \begin{bmatrix} \hat{\mathbf{u}}_{\sigma_1} \\ \hat{\mathbf{u}}_{\sigma_2} \end{bmatrix}.$$

end loop

end loop

loop over all leaf boxes τ

$$\mathbf{q}(I_\tau) = \mathbf{E}_\tau \hat{\mathbf{q}}_\tau + \mathbf{G}_\tau \mathbf{u}(I_\tau).$$

end loop

6.1. A basic compression scheme for leaf nodes. In this section, we describe how to construct for every leaf node τ , interpolatory matrices \mathbf{U}_τ and \mathbf{V}_τ of rank k , and index vectors $\tilde{I}_\tau^{(\text{row})}$ and $\tilde{I}_\tau^{(\text{col})}$ such that, cf. (4.6),

$$(6.1) \quad \mathbf{A}(I_\tau, I_{\tau'}) = \mathbf{U}_\tau \mathbf{A}(\tilde{I}_\tau^{(\text{row})}, \tilde{I}_{\tau'}^{(\text{col})}) \mathbf{V}_{\tau'}^*, \quad \tau \neq \tau'.$$

The first step in the process is to construct for every leaf τ a row of blocks \mathbf{R}_τ and a column of blocks \mathbf{C}_τ via

$$\mathbf{R}_\tau = \mathbf{A}(I_\tau, L_\tau), \quad \text{and} \quad \mathbf{C}_\tau = \mathbf{A}(L_\tau, I_\tau),$$

where L_τ the complement of the index vector I_τ within the full index set,

$$L_\tau = \{1, 2, 3, \dots, N\} \setminus I_\tau.$$

The condition (4.6) implies that \mathbf{R}_τ and \mathbf{C}_τ each have rank at most k . We can therefore construct interpolatory decompositions

$$(6.2) \quad \mathbf{R}_\tau = \mathbf{U}_\tau \mathbf{R}(J_\tau^{(\text{row})}, :),$$

$$(6.3) \quad \mathbf{C}_\tau = \mathbf{C}(:, J_\tau^{(\text{col})}) \mathbf{V}_\tau^*.$$

ALGORITHM 5 (reformatting inverse)

Postprocessing the terms computed Algorithm 3 to obtain an inverse in the standard HBS format.

```

loop over all levels, coarser to finer,  $\ell = 0, 1, 2, \dots, L - 1$ 
  loop over all parent boxes  $\tau$  on level  $\ell$ ,
    Let  $\sigma_1$  and  $\sigma_2$  denote the children of  $\tau$ .
    Define the matrices  $H_{1,1}$ ,  $B_{\sigma_1, \sigma_2}$ ,  $B_{\sigma_2, \sigma_1}$ ,  $H_{2,2}$  so that
       $G_\tau = \begin{bmatrix} H_{1,1} & B_{\sigma_1, \sigma_2} \\ B_{\sigma_2, \sigma_1} & H_{2,2} \end{bmatrix}$ .
       $G_{\sigma_1} \leftarrow G_{\sigma_1} + E_{\sigma_1} H_{1,1} F_{\sigma_1}^*$ .
       $G_{\sigma_2} \leftarrow G_{\sigma_2} + E_{\sigma_2} H_{2,2} F_{\sigma_2}^*$ .
    end loop
  end loop

loop over all leaf boxes  $\tau$ 
   $D_\tau = G_\tau$ .
end loop

```

Now (6.1) holds if we set

$$\tilde{I}_\tau^{(\text{row})} = I_\tau(J_\tau^{(\text{row})}) \quad \text{and} \quad \tilde{I}_\tau^{(\text{col})} = I_\tau(J_\tau^{(\text{col})}).$$

Remark 6.1. It is often useful to use the same basis matrices to span the ranges of both R_τ and C_τ^* . In particular, this can make a substantial difference in the stability of the inversion scheme described in Section 4. To accomplish this, form the matrix $X_\tau = [R_\tau \mid C_\tau^*]$ and then compute an interpolatory factorization

$$(6.4) \quad X_\tau = U_\tau X_\tau(J_\tau, :).$$

Then clearly $R_\tau = U_\tau R(J_\tau, :)$ and $C_\tau = C(:, J_\tau) U_\tau^*$. Enforcing symmetry can slightly increase the HSS-ranks, but typically in a very modest way.

Remark 6.2. In applications, the matrices R_τ and C_τ are typically only *approximately* of rank k and the factorizations (6.2) and (6.3) are then required to hold only to within some preset tolerance ε . Techniques for computing rank-revealing partial factorizations of this type are described in detail in [14, 9].

6.2. An accelerated compression scheme for leaf nodes. We will in this section demonstrate how to rapidly construct matrices U_τ , V_τ and index vectors $J_\tau^{(\text{row})}$, $J_\tau^{(\text{col})}$ such that (6.2) and (6.3) hold. We focus on the construction of U_τ and $J_\tau^{(\text{row})}$ since the construction of V_τ and $J_\tau^{(\text{col})}$ is analogous. While U_τ and $J_\tau^{(\text{row})}$ can in principle be constructed by directly computing an interpolatory decomposition of R_τ this is in practice prohibitively expensive since R_τ is very large. The way to get around this is to exploit analytic properties of the kernel to construct a much smaller matrix $R_\tau^{(\text{small})}$ whose columns span the column space of R_τ . Then we can cheaply form U_τ and $J_\tau^{(\text{row})}$ by factoring this smaller matrix. The process typically over-estimates the rank slightly (since the columns of $R_\tau^{(\text{small})}$ will span a slightly larger space than the columns of R_τ) but this is more than compensated by the dramatic acceleration of the compression step.

To formalize matters, our goal is to construct a small matrix $R_\tau^{(\text{small})}$ such that

$$(6.5) \quad \text{Ran}(R_\tau) \subseteq \text{Ran}(R_\tau^{(\text{small})}).$$




FIGURE 6.1. A contour Γ . (a) Γ_τ is drawn with a bold line. (b) The contour $\Gamma_\tau^{(\text{near})}$ is drawn with a thin solid line and $\Gamma_\tau^{(\text{far})}$ with a dashed line.

Then all we would need to do to compress τ is to construct the interpolatory decomposition

$$(6.6) \quad \mathbf{R}_\tau^{(\text{small})} = \mathbf{U}_\tau \mathbf{R}_\tau^{(\text{small})}(\mathbf{J}_\tau^{(\text{row})}, :)$$

since (6.5) and (6.6) together imply (6.2).

When constructing the matrix $\mathbf{R}_\tau^{(\text{small})}$, we distinguish between near field interaction and far field interactions. The near field interactions cannot readily be compressed, but this is not a problem since they contribute a very small part of \mathbf{R}_τ . To define what we mean by “near” and “far,” we need to introduce some notation. Let Γ_τ denote the segment of Γ associated with the node τ , see Fig. 6.1. We enclose Γ_τ in a circle and then let $\Gamma_\tau^{(\text{proxy})}$ denote a circle with the same center but with a 50% larger radius. We now define $\Gamma_\tau^{(\text{far})}$ as the part of Γ outside of $\Gamma_\tau^{(\text{proxy})}$ and define $\Gamma_\tau^{(\text{near})}$ as the part of Γ inside $\Gamma_\tau^{(\text{proxy})}$ but disjoint from Γ_τ . In other words

$$\Gamma = \Gamma_\tau \cup \Gamma_\tau^{(\text{near})} \cup \Gamma_\tau^{(\text{far})}$$

forms a disjoint partitioning of Γ . We define $L_\tau^{(\text{near})}$ and $L_\tau^{(\text{far})}$ so that

$$\{1, 2, 3, \dots, N\} = I_\tau \cup L_\tau^{(\text{near})} \cup L_\tau^{(\text{far})}$$

forms an analogous disjoint partitioning of the index vector $\{1, 2, \dots, N\}$. We now find that

$$(6.7) \quad \mathbf{R}_\tau = \left[\mathbf{A}(I_\tau, L_\tau^{(\text{near})}) \mid \mathbf{A}(I_\tau, L_\tau^{(\text{far})}) \right] \Pi$$

where Π is a permutation matrix. We will construct a matrix $\mathbf{R}_\tau^{(\text{proxy})}$ such that

$$(6.8) \quad \text{Ran}(\mathbf{A}(I_\tau, L_\tau^{(\text{far})})) \subseteq \text{Ran}(\mathbf{R}_\tau^{(\text{proxy})}),$$

and then we set

$$(6.9) \quad \mathbf{R}_\tau^{(\text{small})} = \left[\mathbf{A}(I_\tau, L_\tau^{(\text{near})}) \mid \mathbf{R}_\tau^{(\text{proxy})} \right].$$

That (6.5) holds is now a consequence of (6.7), (6.8) and (6.9).

All that remains is to construct a small matrix $\mathbf{R}_\tau^{(\text{proxy})}$ such that (6.8) holds. We describe the process for the single layer potential associated with Laplace’s equation (for generalizations, see Remarks 6.4, 6.5, 6.6, and 6.7). Since we use Nyström discretization, the matrix $\mathbf{A}(I_\tau, L_\tau^{(\text{far})})$ in this case represents evaluation on Γ_τ of the harmonic potential generated by a set of point charges on $\Gamma_\tau^{(\text{far})}$. We know from potential theory that any harmonic field generated by charges *outside* $\Gamma_\tau^{(\text{proxy})}$ can be generated by placing an equivalent charge distribution *on* $\Gamma_\tau^{(\text{proxy})}$. Since we only need to capture the field to within some preset precision ε , it is sufficient to place charges on a finite collection $\{\mathbf{z}_j\}_{j=1}^J$ of points on $\Gamma_\tau^{(\text{proxy})}$. In other words, we set

$$\mathbf{R}_\tau^{(\text{proxy})}(i, j) = \log |\mathbf{x}_i - \mathbf{z}_j|, \quad i \in I_\tau, \quad j \in \{1, 2, 3, \dots, J\}.$$

The number of charges J that are needed on the external circle depends on the precision ε required. In fact $J = O(\log(1/\varepsilon))$ as $\varepsilon \rightarrow 0$. We have found that using $J = 50$ points is sufficient to attain ten digits of accuracy or better.

The construction of a small matrix $\mathbf{C}_\tau^{(\text{small})}$ that can be used to construct \mathbf{V}_τ and $\mathbf{J}_\tau^{(\text{col})}$ such that (6.3) holds is entirely analogous since the matrix \mathbf{C}_τ^* is also a map of a charge distribution on $\Gamma_\tau^{(\text{far})}$ to a potential field on Γ_τ . The only caveat is that the rows of \mathbf{C}_τ^* must be scaled by the quadrature weights used in the Nyström method.

Remark 6.3. As an alternative to placing charges on the exterior circle, one could represent the harmonic field generated on Γ_τ by an expansion in the cylindrical harmonic functions $\{r^j \cos(j\theta), r^j \sin(j\theta)\}_{j=0}^J$ where (r, θ) are the polar coordinates of a point in the circle enclosing Γ_τ . The number of functions needed are about the same, but we found it easier to correctly weigh the two terms $\mathbf{A}(I_\tau, L_\tau^{(\text{far})})$ and $\mathbf{R}_\tau^{(\text{proxy})}$ when using proxy charges on the outer circle.

Remark 6.4 (Extension to the double layer kernel). The procedure described directly generalizes to the double layer kernel associated with Laplace's equation. The compression of \mathbf{R}_τ is exactly the same. The compression of \mathbf{C}_τ^* is very similar, but now the target field is the normal derivative of the set of harmonic potentials that can be generated by sources outside $\Gamma_\tau^{(\text{proxy})}$.

Remark 6.5 (Extension to Laplace's equation in \mathbb{R}^3). The scheme described generalizes immediately to BIEs defined on surfaces in \mathbb{R}^3 . The circles must be replaced by spheres, and the complexity is no longer linear, but the method remains competitive at low and intermediate accuracies [12].

Remark 6.6 (Extension to Helmholtz and other equations). The scheme described has been generalized to the single and double layer potentials associated with Helmholtz equation, see [21]. The only complication happens the frequency is close to a resonant frequency of the proxy circle. This potential problem is avoided by placing both monopoles and dipoles on the proxy circle.

Remark 6.7 (Extension to integral equations on the line). The acceleration gets particularly simple for integral equations on a line with smooth non-oscillatory kernels. In this case, the range of $\mathbf{A}(I_\tau, L_\tau^{(\text{far})})$ can typically be represented by a short expansion in a generic set of basis functions such as, e.g., Legendre polynomials.

6.3. Compression of higher levels. The method described in Section 6.2 rapidly constructs the matrices \mathbf{U}_τ , \mathbf{V}_τ and the index vector $\mathbf{J}_\tau^{(\text{row})}$, $\mathbf{J}_\tau^{(\text{col})}$ for any leaf node τ . Using the notation introduced in Section 4.4, it computes the matrices $\underline{\mathbf{U}}^{(L)}$, $\underline{\mathbf{V}}^{(L)}$, and $\tilde{\mathbf{A}}^{(L)}$. It is important to note that when the interpolatory decomposition is used, $\tilde{\mathbf{A}}^{(L)}$ is in fact a submatrix of \mathbf{A} , and is represented *implicitly* by specifying the relevant index vectors. To be precise, if τ and τ' are two nodes on level $L - 1$, with children σ_1, σ_2 and σ'_1, σ'_2 , respectively, then

$$\mathbf{A}_{\tau, \tau'} = \begin{bmatrix} \mathbf{A}(\tilde{I}_{\sigma_1}^{(\text{row})}, \tilde{I}_{\sigma'_1}^{(\text{col})}) & \mathbf{A}(\tilde{I}_{\sigma_1}^{(\text{row})}, \tilde{I}_{\sigma'_2}^{(\text{col})}) \\ \mathbf{A}(\tilde{I}_{\sigma_2}^{(\text{row})}, \tilde{I}_{\sigma'_1}^{(\text{col})}) & \mathbf{A}(\tilde{I}_{\sigma_2}^{(\text{row})}, \tilde{I}_{\sigma'_2}^{(\text{col})}) \end{bmatrix}.$$

The observation that $\tilde{\mathbf{A}}^{(L)}$ is a submatrix of \mathbf{A} is critical. It implies that the matrices $\underline{\mathbf{U}}^{(L-1)}$, $\underline{\mathbf{V}}^{(L-1)}$, and $\tilde{\mathbf{A}}^{(L-1)}$ can be computed using the strategy of Section 6.2 *without any modifications*.

6.4. Approximation errors. As mentioned in Remark 6.2, factorizations such as (6.2), (6.3), (6.4), (6.6) are in practice only required to hold to within some preset tolerance. In a single-level scheme, it is straight-forward to choose a local tolerance in such a way that $\|\mathbf{A} - \mathbf{A}_{\text{approx}}\| \leq \varepsilon$ holds to within some given global tolerance ε . In a multi-level scheme, it is rather difficult to

predict how errors aggregate across levels, in particular when the basis matrices \mathbf{U}_τ and \mathbf{V}_τ are not orthonormal. As an empirical observation, we have found that such error propagation is typically very mild and the error $\|\mathbf{A} - \mathbf{A}_{\text{approx}}\|$ is very well predicted by the local tolerance in the interpolatory decomposition.

While there are as of yet no *a priori* error guarantees, it is often possible to produce highly accurate *a posteriori* error estimates. To this end, let $\mathbf{q} = \mathbf{A}^{-1} \mathbf{f}$ and $\mathbf{q}_{\text{approx}} = \mathbf{A}_{\text{approx}}^{-1} \mathbf{f}$ denote the exact and the approximate solution, respectively. Then

$$\|\mathbf{q}_{\text{approx}} - \mathbf{q}\| = \|\mathbf{A}_{\text{approx}}^{-1} \mathbf{A} \mathbf{q} - \mathbf{A}_{\text{approx}}^{-1} \mathbf{A}_{\text{approx}} \mathbf{q}\| \leq \|\mathbf{A}_{\text{approx}}^{-1}\| \|\mathbf{A} - \mathbf{A}_{\text{approx}}\| \|\mathbf{q}\|.$$

We can very rapidly and accurately estimate $\|\mathbf{A}_{\text{approx}}^{-1}\|$ via a power iteration since we have access to a fast matrix-vector multiply for $\mathbf{A}_{\text{approx}}^{-1}$. The factor $\|\mathbf{A} - \mathbf{A}_{\text{approx}}\|$ can similarly be estimated whenever we have access to an independent fast matrix-vector multiplication for \mathbf{A} . For all BIEs discussed in this paper, the Fast Multipole Method can serve this function. If it is found that the factor $\|\mathbf{A}_{\text{approx}}^{-1}\| \|\mathbf{A} - \mathbf{A}_{\text{approx}}\|$ is larger than desired, the compression can be rerun with a smaller local tolerance. (The argument in this section assumes that the inversion is truly exact; a careful error analysis must also account for propagation of round-off errors.)

7. NUMERICAL EXAMPLES

The performance of the direct solver was tested on linear systems arising upon the Nyström discretization of the BIE

$$(7.1) \quad \frac{1}{2} q(\mathbf{x}) + \int_{\Gamma} \frac{\mathbf{n}(\mathbf{x}') \cdot (\mathbf{x} - \mathbf{x}')}{2\pi |\mathbf{x} - \mathbf{x}'|^2} q(\mathbf{x}') dl(\mathbf{x}') = f(\mathbf{x}), \quad \mathbf{x} \in \Gamma,$$

where $\mathbf{n}(\mathbf{y})$ is a unit normal to the contour Γ . Equation (7.1) is a standard BIE representation of the Laplace equation on a domain bordered by Γ when Dirichlet boundary data is specified [19, Sec. 6.3]. Three geometries were considered:

- **Smooth star:** The star domain illustrated in Figure 7.1(a) is discretized via Gaussian quadrature as described in Section 2.3. In the experiment, we fix the size of the computational domain and increase the number of discretization points. This problem is artificial in the sense that the largest experiments use *far* more quadrature points than what is required for any reasonable level of accuracy. It was included simply to demonstrate the asymptotic scaling of the method.
- **Star with corners:** The contour Γ consists of ten segments of circles with different radii as illustrated in Figure 7.1(b). We started with a discretization with 6 panels per segment and 17 Gaussian quadrature nodes per panel. Then grid refinement as described in [17, 5] (with a so called “simply graded mesh”) was used to increase the number of discretization points, cf. Remark 2.2.
- **Snake:** The contour Γ consists of two sine waves with amplitude 1 that are vertically separated by a distance of 0.2, see Figure 7.1(c). At the end points, two vertical straight lines connect the waves. Composite Gaussian quadrature was used with 25 nodes per panel, four panels on each vertical straight line (refined at the corners to achieve 10 digits of accuracy), and then 10 panels were used per wave-length. The width of the contour was increase from 2 full periods of the wave to 200, and the number of discretization points N was increased accordingly.

A compressed representation of the coefficient matrix was for each example computed via Matlab implementations of the compression scheme described in Section 6. Then Fortran 77 implementations of Algorithms 3 (inversion of an HBS matrix), 5 (conversion to standard HBS format), and 2 (matrix-vector multiply) were used to directly solve the respective linear systems. All codes were executed on a desktop computer with 2.8GHz Intel i7 processor and 2GB of




FIGURE 7.1. The contours Γ used in the numerical experiments in Section 7.
 (a) Smooth star. (b) Star with corners. (c) Snake.

RAM. The speed of the compression step can be improved significantly by moving to a Fortran implementation, but since this step is somewhat idiosyncratic to each specific problem, we believe that is representative to report the times of an unoptimized Matlab code.

For all experiments, the local tolerance in the compression step was set to $\epsilon = 10^{-10}$. (In other words, the interpolatory factorization in (6.6) was required to produce a residual in the Frobenius norm no larger than ϵ .)

To assess the accuracy of the direct solver, the error $\|\mathbf{A} - \mathbf{A}_{\text{approx}}\|$ was estimated via a power iteration after each compression (the matrix \mathbf{A} and its transpose were applied via the classical FMM [13] run at very high accuracy). The quantity $\|\mathbf{A}_{\text{approx}}^{-1}\|$ was also estimated. The results are reported in Figure 7.3. The quantities reported bound the overall error in the direct solver, see 6.4.

8. EXTENSIONS AND FUTURE WORK

While the numerical examples in this paper concerned only BIEs associated with Laplace's equation, then method can readily be extended to other elliptic problems, as long as the kernel is not too oscillatory. The extension to Helmholtz equation was reported in [21] and [22].

The direct solver described can also be applied to BIEs defined on surfaces in \mathbb{R}^3 without almost any modifications [12]. The complexity then grows from $O(N)$ to $O(N^{1.5})$ for the inversion step, while $O(N)$ complexity is retained for applying the computed inverse. The complexity of the inversion step can be improved by doing a “recursion on dimension” in a manner similar to the recently described $O(N)$ nested dissection schemes reported in [25, 7, 11]. This work is currently in progress.

The error analysis in the present paper is very rudimentary. Numerical experiments indicate that the method is stable and accurate, but this has not yet been demonstrated in any case other than that of symmetric positive definite matrices.

Acknowledgements: The work reported was supported by NSF grants DMS0748488 and DMS0941476.

REFERENCES

- [1] K. E. Atkinson, *The numerical solution of integral equations of the second kind*, Cambridge University Press, Cambridge, 1997. MR 99d:65364
- [2] J. Barnes and P. Hut, *A hierarchical $o(n \log n)$ force-calculation algorithm*, Nature **324** (1986), no. 4.
- [3] G. Beylkin, R. Coifman, and V. Rokhlin, *Wavelets in numerical analysis*, Wavelets and their applications, Jones and Bartlett, Boston, MA, 1992, pp. 181–210.
- [4] S. Börm, *Efficient numerical methods for non-local operators: \mathcal{H}^2 -matrix compression, algorithms and analysis*, European Mathematics Society, 2010.

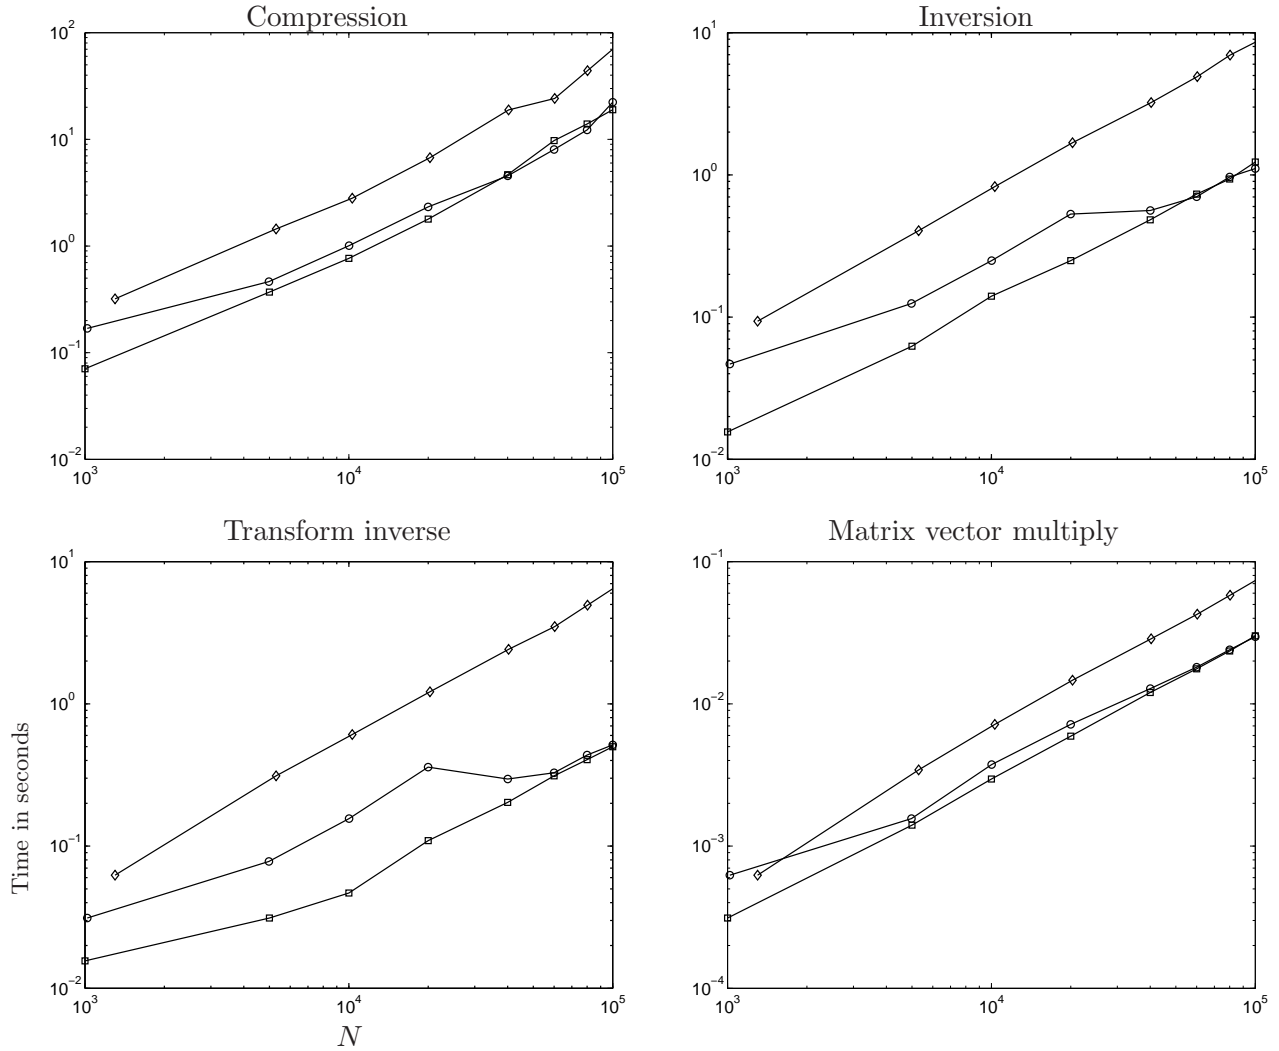


FIGURE 7.2. The four graphs give the times required for the four steps in the direct solver. Within each graph, the three lines correspond to the three different contours considered: \square – Smooth star, \circ – Star with corners, \diamond – Snake.

- [5] J. Bremer and V. Rokhlin, *Efficient discretization of Laplace boundary integral equations on polygonal domains*, J. Comput. Phys. **229** (2010), 2507–2525.
- [6] S. Chandrasekaran and M. Gu, *A divide-and-conquer algorithm for the eigendecomposition of symmetric block-diagonal plus semiseparable matrices*, Numer. Math. **96** (2004), no. 4, 723–731.
- [7] S. Chandrasekaran, M. Gu, X.S. Li, and J. Xia, *Superfast multifrontal method for large structured linear systems of equations*, SIAM J. Matrix Anal. Appl. **31** (2009), 1382–1411.
- [8] S. Chandrasekaran, M. Gu, X.S. Li, and J. Xia, *Fast algorithms for hierarchically semiseparable matrices*, Numer. Linear Algebra Appl. **17** (2010), 953–976.
- [9] H. Cheng, Z. Gimbutas, P.G. Martinsson, and V. Rokhlin, *On the compression of low rank matrices*, SIAM Journal of Scientific Computing **26** (2005), no. 4, 1389–1404.
- [10] Gene H. Golub and Charles F. Van Loan, *Matrix computations*, third ed., Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, 1996.
- [11] L. Grasedyck, R. Kriemann, and S. Le Borne, *Domain decomposition based h-lu preconditioning*, Numerische Mathematik **112** (2009), 565–600.




FIGURE 7.3. Error information for each of the domains: \square – smooth star, \circ – star with corners, \diamond – snake.

- [12] L. Greengard, D. Gueyffier, P.G. Martinsson, and V. Rokhlin, *Fast direct solvers for integral equations in complex three-dimensional domains*, Acta Numerica **18** (2009), 243–275.
- [13] L. Greengard and V. Rokhlin, *A fast algorithm for particle simulations*, J. Comput. Phys. **73** (1987), no. 2, 325–348. MR MR918448 (88k:82007)
- [14] Ming Gu and Stanley C. Eisenstat, *Efficient algorithms for computing a strong rank-revealing QR factorization*, SIAM J. Sci. Comput. **17** (1996), no. 4, 848–869. MR 97h:65053
- [15] W. Hackbusch, *The panel clustering technique for the boundary element method (invited contribution)*, Boundary elements IX, Vol. 1 (Stuttgart, 1987), Comput. Mech., Southampton, 1987, pp. 463–474.
- [16] Wolfgang Hackbusch, *A sparse matrix arithmetic based on H-matrices; Part I: Introduction to H-matrices*, Computing **62** (1999), 89–108.
- [17] J. Helsing and R. Ojala, *Corner singularities for elliptic problems: Integral equations, graded meshes, quadrature, and compressed inverse preconditioning*, J. Comput. Phys. **227** (2008), 8820–8840.
- [18] S. Kapur and V. Rokhlin, *High-order corrected trapezoidal quadrature rules for singular functions*, SIAM J. Numer. Anal. **34** (1997), 1331–1356.

- [19] Rainer Kress, *Linear integral equations*, second ed., Applied Mathematical Sciences, vol. 82, Springer-Verlag, New York, 1999.
- [20] P.G. Martinsson, *A fast algorithm for compressing a matrix into a data-sparse format via randomized sampling*, Tech. report, 2008, arXiv.org report 0806.2339.
- [21] P.G. Martinsson and V. Rokhlin, *A fast direct solver for boundary integral equations in two dimensions*, J. Comp. Phys. **205** (2005), no. 1, 1–23.
- [22] ———, *A fast direct solver for scattering problems involving elongated structures*, Journal of Computational Physics **221** (2007), 288 – 302.
- [23] E. Michielssen, A. Boag, and W. C. Chew, *Scattering from elongated objects: direct solution in $O(N \log^2 N)$ operations*, IEE Proc. Microw. Antennas Propag. **143** (1996), no. 4, 277 – 283.
- [24] S. T. O’Donnell and V. Rokhlin, *A fast algorithm for the numerical evaluation of conformal mappings*, SIAM J. Sci. Stat. Comput. **10** (1989), 475–487.
- [25] P. Schmitz and L. Ying, *A fast direct solver for elliptic problems on general meshes in 2d*, 2010, In review.
- [26] Z. Sheng, P. Dewilde, and S. Chandrasekaran, *Algorithms to solve hierarchically semi-separable systems*, System theory, the Schur algorithm and multidimensional analysis, Oper. Theory Adv. Appl., vol. 176, Birkhäuser, Basel, 2007, pp. 255–294.
- [27] Page Starr and Vladimir Rokhlin, *On the numerical solution of two-point boundary value problems. II*, Comm. Pure Appl. Math. **47** (1994), no. 8, 1117–1159.
- [28] H. Xiao, V. Rokhlin, and N. Yarvin, *Prolate spheroidal wavefunctions, quadrature and interpolation*, Inverse Problems **17** (2001), no. 4.
- [29] P. Young and P.G. Martinsson, *A direct solver for the rapid solution of boundary integral equations on axisymmetric surfaces in three dimensions*, 2009, arxiv.org report #1002.2001.